Behavioral Economics

Topic 1: Choice under Uncertainty

Lecture 2: The Standard Model — Expected Utility

Professor Bushong

MICHIGAN STATE UNIVERSITY

Taking Stock

Last class:

- Introduction to the course and material
- Any questions after reading the syllabus?

Taking Stock

Last class:

- Introduction to the course and material
- Any questions after reading the syllabus?

Today:

- Start on our first major topic: choice under uncertainty
- Today we'll focus on the standard model

Taking Stock

Last class:

- Introduction to the course and material
- Any questions after reading the syllabus?

Today:

- Start on our first major topic: choice under uncertainty
- Today we'll focus on the standard model

Logistics:

- Problem Set 0 is posted – optional "background" material on optimization
- Problem Set 1 is posted – the first real problem set; due Tuesday, September 15
    - Tip: start early – don't wait until the last minute

Our Outline

A Historical Perspective

The study of risk is perhaps the oldest element of “decision science” or behavioral economics.

A Historical Perspective

The study of risk is perhaps the oldest element of “decision science” or behavioral economics.

Renaissance scientists and mathematicians mused about probability and concocted puzzles around games of chance.

A Historical Perspective

The study of risk is perhaps the oldest element of “decision science” or behavioral economics.

Renaissance scientists and mathematicians mused about probability and concocted puzzles around games of chance.

- 1494: Franciscan monk Luca Pacioli proposed “the problem of points"—which asks how one should divide the stakes in an incomplete game.

A Historical Perspective

The study of risk is perhaps the oldest element of “decision science” or behavioral economics.

Renaissance scientists and mathematicians mused about probability and concocted puzzles around games of chance.

- 1494: Franciscan monk Luca Pacioli proposed “the problem of points"—which asks how one should divide the stakes in an incomplete game.
- 1654: French mathematicians Blaise Pascal and Pierre de Fermat developed a way to determine the likelihood of each possible result of a simple game (balla).

A Historical Perspective

The study of risk is perhaps the oldest element of "decision science" or behavioral economics.

Renaissance scientists and mathematicians mused about probability and concocted puzzles around games of chance.

- 1494: Franciscan monk Luca Pacioli proposed "the problem of points"—which asks how one should divide the stakes in an incomplete game.
- 1654: French mathematicians Blaise Pascal and Pierre de Fermat developed a way to determine the likelihood of each possible result of a simple game (balla).

Today's lecture will cover insights from the mid 18th century to modern approaches.

Defining the Environment

A key part of good economic modeling: precisely define your environment and terms

Defining the Environment

A key part of good economic modeling: precisely define your environment and terms

My view: it is impossible to operate without precise definitions of things.

- E.g., if we don't define what it means for someone to "get sick", it's hard to have productive conversations about appropriate health policy.

Defining the Environment

A key part of good economic modeling: precisely define your environment and terms

My view: it is impossible to operate without precise definitions of things.

- E.g., if we don't define what it means for someone to "get sick", it's hard to have productive conversations about appropriate health policy.

This seems to be especially important when we discuss "intuitive" things like "risk".

- Your task: understand the definitions in this course — they are *precise*.

Defining the Environment

Definition: A *lottery* (or *gamble* or *risky prospect*) is a set of possible outcomes and a probability of each outcome occurring.

Defining the Environment

Definition: A *lottery* (or *gamble* or *risky prospect*) is a set of possible outcomes and a probability of each outcome occurring.

Example A: If you go to a roulette table and place \$10 on **black**, you have just purchased a lottery. With probability $18/38$ you will receive \$20, and with probability $20/38$ you will receive \$0.

Defining the Environment

Definition: A *lottery* (or *gamble* or *risky prospect*) is a set of possible outcomes and a probability of each outcome occurring.

Example A: If you go to a roulette table and place \$10 on **black**, you have just purchased a lottery. With probability $18/38$ you will receive \$20, and with probability $20/38$ you will receive \$0.

Example B: When you buy a new car, you are buying a lottery. For instance, it might be that with probability $1/2$ it will be a car that you love (high value), with probability $2/5$ it will be a car that only adequately serves your needs (low value), and with probability $1/10$ it will be a “lemon" (worthless).

Defining the Environment

Definition: A *lottery* (or *gamble* or *risky prospect*) is a set of possible outcomes and a probability of each outcome occurring.

Example A: If you go to a roulette table and place \$10 on **black**, you have just purchased a lottery. With probability $18/38$ you will receive \$20, and with probability $20/38$ you will receive \$0.

Example B: When you buy a new car, you are buying a lottery. For instance, it might be that with probability $1/2$ it will be a car that you love (high value), with probability $2/5$ it will be a car that only adequately serves your needs (low value), and with probability $1/10$ it will be a “lemon" (worthless).

Example C: More abstractly, you might get x_1 with probability p_1, x_2 with probability p_2, and x_3 with probability p_3 (where $p_1 + p_2 + p_3 = 1$).

Defining the Environment

- We can write lotteries as a probability tree.

Defining the Environment

- We can write lotteries as a probability tree.
- Or as a vector of outcome-probability pairs:

Defining the Environment

- We can write lotteries as a probability tree.
- Or as a vector of outcome-probability pairs:

  Example A: Payoff $= (\ \$20, 18/38\ ;\ \$0, 20/38\ )$.

Defining the Environment

- We can write lotteries as a probability tree.
- Or as a vector of outcome-probability pairs:

  Example A: Payoff $= (\ \$20, 18/38\ ;\ \$0, 20/38\ )$.

  Example B: Car Value $= (\ \text{high}, 1/2\ ;\ \text{low}, 2/5\ ;\ \text{worthless}, 1/10)$.

Defining the Environment

- We can write lotteries as a probability tree.
- Or as a vector of outcome-probability pairs:

  Example A: Payoff $= (\ \$20, 18/38\ ;\ \$0, 20/38\ )$.

  Example B: Car Value $= (\ \text{high}, 1/2\ ;\ \text{low}, 2/5\ ;\ \text{worthless}, 1/10)$.
- Or as a type of equation:

$$\text{Example A:} \qquad \text{Payoff} = \begin{cases} \$20 & \text{with prob } 18/38 \\ \$0 & \text{with prob } 20/38 \end{cases}$$

Defining the Environment

- We can write lotteries as a probability tree.
- Or as a vector of outcome-probability pairs:

  Example A: Payoff $= (\ \$20, 18/38\ ;\ \$0, 20/38\ )$.

  Example B: Car Value $= (\ \text{high}, 1/2\ ;\ \text{low}, 2/5\ ;\ \text{worthless}, 1/10)$.

- Or as a type of equation:

$$\text{Example A:} \qquad \text{Payoff} = \begin{cases} \$20 & \text{with prob } 18/38 \\ \$0 & \text{with prob } 20/38 \end{cases}$$

$$\text{Example B:} \qquad \text{Car Value} = \begin{cases} \text{high value} & \text{with prob } 1/2 \\ \text{low value} & \text{with prob } 2/5 \\ \text{worthless} & \text{with prob } 1/10 \end{cases}$$

Defining the Environment

In the realm of choice under uncertainty, we examine people's choices of lotteries.

Defining the Environment

In the realm of choice under uncertainty, we examine people's choices of lotteries.

Next step: Develop models of behavior.

Defining the Environment

In the realm of choice under uncertainty, we examine people's choices of lotteries.

Next step: Develop models of behavior.

- Suppose you face a choice between two lotteries. How do you decide?

Model #1: Expected Value Theory (EV)

One model of behavior: People choose the option with the largest “expected value”.

Model #1: Expected Value Theory (EV)

One model of behavior: People choose the option with the largest "expected value".

Definition: The *expected value* of a lottery $\mathbf{x} = (\ x_1, p_1\ ;\ ...\ ;\ x_n, p_n\ )$ is

$$EV(\mathbf{x}) \equiv \sum_{i=1}^{n} p_i\, x_i.$$

Model #1: Expected Value Theory (EV)

One model of behavior: People choose the option with the largest "expected value".

Definition: The *expected value* of a lottery $\mathbf{x} = (\ x_1, p_1\ ;\ ...\ ;\ x_n, p_n\ )$ is

$$EV(\mathbf{x}) \equiv \sum_{i=1}^{n} p_i\ x_i.$$

E.g. the expected value of a 50/50 gamble in which you might win \$10 and otherwise win nothing is \$5.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)

Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get $2.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.
- If I flip two TAILS and then a HEADS, you get \$8.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.
- If I flip two TAILS and then a HEADS, you get \$8.
- If I flip three TAILS and then a HEADS, you get \$16.

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get $2.
- If I flip one TAILS and then a HEADS, you get $4.
- If I flip two TAILS and then a HEADS, you get $8.
- If I flip three TAILS and then a HEADS, you get $16.
- And so forth....

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)

Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.
- If I flip two TAILS and then a HEADS, you get \$8.
- If I flip three TAILS and then a HEADS, you get \$16.
- And so forth....

How much are you willing to pay for this bet?

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)

Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.
- If I flip two TAILS and then a HEADS, you get \$8.
- If I flip three TAILS and then a HEADS, you get \$16.
- And so forth....

How much are you willing to pay for this bet?

How much would an expected-value maximizer pay?

A Problem for EV

St. Petersburg Paradox (posed by Nicolas Bernoulli, 1713)
Consider the following bet: I'm going to flip a coin, and I'm going to keep on flipping it until I flip a HEADS. Then you'll be paid as a function of how many times we flip. Specifically:

- If I immediately flip a HEADS, you get \$2.
- If I flip one TAILS and then a HEADS, you get \$4.
- If I flip two TAILS and then a HEADS, you get \$8.
- If I flip three TAILS and then a HEADS, you get \$16.
- And so forth....

How much are you willing to pay for this bet?

How much would an expected-value maximizer pay? The "paradox":

- The *EV* of this bet is ∞, but people are unwilling to pay much for it. Hence, *EV* is not a good description of people's choices.

St. Petersburg Paradox

"Now it is highly probable that any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed." – Daniel Bernoulli, 1738 (translated in 1954)

St. Petersburg Paradox

"Now it is highly probable that any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed." – Daniel Bernoulli, 1738 (translated in 1954)

Bernoulli's observation:

- People have **diminishing marginal utility** of money.

St. Petersburg Paradox

"Now it is highly probable that any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed." – Daniel Bernoulli, 1738 (translated in 1954)

Bernoulli's observation:

- People have **diminishing marginal utility** of money.
- Just because the gamble is worth infinite money does not mean it is infinitely valuable to a would-be gambler.

St. Petersburg Paradox

"Now it is highly probable that any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed." – Daniel Bernoulli, 1738 (translated in 1954)

Bernoulli's observation:

- People have **diminishing marginal utility** of money.
- Just because the gamble is worth infinite money does not mean it is infinitely valuable to a would-be gambler.

Immediate implication: People should maximize **utility**, not money.

St. Petersburg Paradox

"Now it is highly probable that any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed." – Daniel Bernoulli, 1738 (translated in 1954)

Bernoulli's observation:

- People have **diminishing marginal utility** of money.
- Just because the gamble is worth infinite money does not mean it is infinitely valuable to a would-be gambler.

Immediate implication: People should maximize **utility**, not money.

Implication 2: Maximizing utility implies people turn down many risks.

What's Missing from EV? "Risk Aversion"

Definition: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

What's Missing from EV? "Risk Aversion"

Definition: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

Definition: A person is *risk-neutral* if for any (risky) lottery $\mathbf{x}$ she is indifferent between $EV(\mathbf{x})$ for sure vs. lottery $\mathbf{x}$.

What's Missing from EV? "Risk Aversion"

Definition: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

Definition: A person is *risk-neutral* if for any (risky) lottery $\mathbf{x}$ she is indifferent between $EV(\mathbf{x})$ for sure vs. lottery $\mathbf{x}$.

Definition: A person is *risk-loving* if for any (risky) lottery $\mathbf{x}$ she prefers to have lottery $\mathbf{x}$ instead of $EV(\mathbf{x})$ for sure.

What's Missing from EV? "Risk Aversion"

<u>Definition</u>: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

<u>Definition</u>: A person is *risk-neutral* if for any (risky) lottery $\mathbf{x}$ she is indifferent between $EV(\mathbf{x})$ for sure vs. lottery $\mathbf{x}$.

<u>Definition</u>: A person is *risk-loving* if for any (risky) lottery $\mathbf{x}$ she prefers to have lottery $\mathbf{x}$ instead of $EV(\mathbf{x})$ for sure.

- Evidence suggests people are risk-averse.

What's Missing from EV? "Risk Aversion"

Definition: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

Definition: A person is *risk-neutral* if for any (risky) lottery $\mathbf{x}$ she is indifferent between $EV(\mathbf{x})$ for sure vs. lottery $\mathbf{x}$.

Definition: A person is *risk-loving* if for any (risky) lottery $\mathbf{x}$ she prefers to have lottery $\mathbf{x}$ instead of $EV(\mathbf{x})$ for sure.

- Evidence suggests people are risk-averse.
- Also your intuition suggests people are risk-averse.

What's Missing from EV? "Risk Aversion"

<u>**Definition**</u>: A person is *risk-averse* if for any (risky) lottery $\mathbf{x}$ she prefers to have $EV(\mathbf{x})$ for sure instead of lottery $\mathbf{x}$.

<u>**Definition**</u>: A person is *risk-neutral* if for any (risky) lottery $\mathbf{x}$ she is indifferent between $EV(\mathbf{x})$ for sure vs. lottery $\mathbf{x}$.

<u>**Definition**</u>: A person is *risk-loving* if for any (risky) lottery $\mathbf{x}$ she prefers to have lottery $\mathbf{x}$ instead of $EV(\mathbf{x})$ for sure.

- Evidence suggests people are risk-averse.
- Also your intuition suggests people are risk-averse.
- Except it's not true all the time . . . we'll discuss this later.

Model #2: Expected Utility Theory (EU)

A second model of behavior, building on economic models and Bernoulli's intuition: People choose the option with the largest **expected utility**.

Model #2: Expected Utility Theory (EU)

A second model of behavior, building on economic models and Bernoulli's intuition: People choose the option with the largest **expected utility**.

Definition: The *expected utility* of a lottery $\mathbf{x} = (x_1, p_1; ...; x_n, p_n)$ is

$$EU(\mathbf{x}) \equiv \sum_{i=1}^{n} p_i \, u(x_i).$$

Model #2: Expected Utility Theory (EU)

A second model of behavior, building on economic models and Bernoulli's intuition: People choose the option with the largest **expected utility**.

<u>**Definition**</u>: The *expected utility* of a lottery $\mathbf{x} = (x_1, p_1; ...; x_n, p_n)$ is

$$EU(\mathbf{x}) \equiv \sum_{i=1}^{n} p_i \, u(x_i).$$

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 50, \tfrac{1}{2})?$$

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 50, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (200, \tfrac{1}{2}; 0, \tfrac{1}{2})?$$

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 50, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (200, \tfrac{1}{2}; 0, \tfrac{1}{2})?$$

Point: If we put no restrictions on the utility function $u(x)$, then expected-utility theory can explain **any** individual choice

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 50, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (200, \tfrac{1}{2}; 0, \tfrac{1}{2})?$$

Point: If we put no restrictions on the utility function $u(x)$, then expected-utility theory can explain **any** individual choice (e.g., in examples above, what if $u(x) = -x$ vs $u(x) = x$?)

Basic Implications of EU

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 100, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (100, \tfrac{1}{2}; 50, \tfrac{1}{2})?$$

Could an expected-utility maximizer (EU maximizer) choose

$$(100, \tfrac{1}{2}; 0, \tfrac{1}{2}) \text{ over } (200, \tfrac{1}{2}; 0, \tfrac{1}{2})?$$

Point: If we put no restrictions on the utility function $u(x)$, then expected-utility theory can explain **any** individual choice (e.g., in examples above, what if $u(x) = -x$ vs $u(x) = x$?)

- And we're left with a pretty unsatisfying theory of behavior.

Basic Implications of EU

Three Comments:

(1) We say that expected-utility theory can "explain" a choice if there exists a utility function $u(x)$ such that an EU maximizer with utility function $u(x)$ would make that choice.

(2) Even when we put no restrictions on the utility function $u(x)$, although an individual choice cannot violate EU, combinations of choices can violate EU.

- E.g., $(100, \frac{1}{2}; 0, \frac{1}{2}) \succ (200, \frac{1}{2}; 0, \frac{1}{2})$

  AND $(100, \frac{1}{3}; 500, \frac{2}{3}) \prec (200, \frac{1}{3}; 500, \frac{2}{3})$

(3) When we put restrictions on the utility function $u(x)$, then even an individual choice could violate EU.

Basic Implications of EU

Natural restriction: We typically assume that **"more is better"**, meaning u is increasing in x.

Basic Implications of EU

Natural restriction: We typically assume that **"more is better"**, meaning u is increasing in x.

Definition: Lottery **x** *dominates* lottery **y** (or equivalently, lottery **y** *is dominated by* lottery **x**) if for every amount z the probability of getting at least z in lottery **x** is at least as large as in lottery **y** — and strictly larger for at least one z.

Basic Implications of EU

Natural restriction: We typically assume that **"more is better"**, meaning u is increasing in x.

Definition: Lottery **x** *dominates* lottery **y** (or equivalently, lottery **y** *is dominated by* lottery **x**) if for every amount z the probability of getting at least z in lottery **x** is at least as large as in lottery **y** — and strictly larger for at least one z.

Examples:

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \qquad \text{vs.} \qquad (100, \tfrac{1}{2}; 0, \tfrac{1}{2})$$

Basic Implications of EU

Natural restriction: We typically assume that "**more is better**", meaning u is increasing in x.

Definition: Lottery **x** *dominates* lottery **y** (or equivalently, lottery **y** *is dominated by* lottery **x**) if for every amount z the probability of getting at least z in lottery **x** is at least as large as in lottery **y** — and strictly larger for at least one z.

Examples:

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \qquad \text{vs.} \qquad (100, \tfrac{1}{2}; 0, \tfrac{1}{2})$$

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \qquad \text{vs.} \qquad (150, \tfrac{1}{2}; 50, \tfrac{1}{2})$$

Basic Implications of EU

Natural restriction: We typically assume that **"more is better"**, meaning u is increasing in x.

Definition: Lottery **x** *dominates* lottery **y** (or equivalently, lottery **y** *is dominated by* lottery **x**) if for every amount z the probability of getting at least z in lottery **x** is at least as large as in lottery **y** — and strictly larger for at least one z.

Examples:

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \quad \text{vs.} \quad (100, \tfrac{1}{2}; 0, \tfrac{1}{2})$$

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \quad \text{vs.} \quad (150, \tfrac{1}{2}; 50, \tfrac{1}{2})$$

$$(200, \tfrac{1}{3}; 150, \tfrac{1}{3}, 75, \tfrac{1}{3}) \quad \text{vs.} \quad (150, \tfrac{1}{2}; 75, \tfrac{1}{2})$$

Basic Implications of EU

Natural restriction: We typically assume that **"more is better"**, meaning u is increasing in x.

Definition: Lottery **x** *dominates* lottery **y** (or equivalently, lottery **y** *is dominated by* lottery **x**) if for every amount z the probability of getting at least z in lottery **x** is at least as large as in lottery **y** — and strictly larger for at least one z.

Examples:

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \quad \text{vs.} \quad (100, \tfrac{1}{2}; 0, \tfrac{1}{2})$$

$$(200, \tfrac{1}{2}; 0, \tfrac{1}{2}) \quad \text{vs.} \quad (150, \tfrac{1}{2}; 50, \tfrac{1}{2})$$

$$(200, \tfrac{1}{3}; 150, \tfrac{1}{3}, 75, \tfrac{1}{3}) \quad \text{vs.} \quad (150, \tfrac{1}{2}; 75, \tfrac{1}{2})$$

Result: Under EU with "more is better," a person will never choose a dominated lottery.

Basic Implications of EU

To incorporate risk-aversion (or risk-seeking), we must assume more about the utility function.

Basic Implications of EU

To incorporate risk-aversion (or risk-seeking), we must assume more about the utility function.

Result: Under expected-utility theory, a person is risk-averse if and only if her utility function is concave ($u''(x) < 0$ — slope is getting flatter as x gets bigger)

Basic Implications of EU

To incorporate risk-aversion (or risk-seeking), we must assume more about the utility function.

Result: Under expected-utility theory, a person is risk-averse if and only if her utility function is concave ($u''(x) < 0$ — slope is getting flatter as x gets bigger)

Result: Under expected-utility theory, a person is risk-neutral if and only if her utility function is linear ($u''(x) = 0$ — slope is constant)

Basic Implications of EU

To incorporate risk-aversion (or risk-seeking), we must assume more about the utility function.

Result: Under expected-utility theory, a person is risk-averse if and only if her utility function is concave ($u''(x) < 0$ — slope is getting flatter as x gets bigger)

Result: Under expected-utility theory, a person is risk-neutral if and only if her utility function is linear ($u''(x) = 0$ — slope is constant)

Result: Under expected-utility theory, a person is risk-loving if and only if her utility function is convex($u''(x) > 0$ — slope is getting steeper as x gets bigger)

Basic Implications of EU

To incorporate risk-aversion (or risk-seeking), we must assume more about the utility function.

Result: Under expected-utility theory, a person is risk-averse if and only if her utility function is concave ($u''(x) < 0$ — slope is getting flatter as x gets bigger)

Result: Under expected-utility theory, a person is risk-neutral if and only if her utility function is linear ($u''(x) = 0$ — slope is constant)

Result: Under expected-utility theory, a person is risk-loving if and only if her utility function is convex($u''(x) > 0$ — slope is getting steeper as x gets bigger)

Note: In this class, whenever we say that someone is a risk-averse (or risk-neutral or risk-loving) expected-utility maximizer, we'll implicitly assume more is better as well.

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Proper way to use expected-utility theory:

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Proper way to use expected-utility theory:

Let w be a person's *wealth*.

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Proper way to use expected-utility theory:

Let w be a person's *wealth*.

Let $\mathbf{x} = (x_1, p_1; ...; x_n, p_n)$ be a lottery (or gamble or risky prospect).

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Proper way to use expected-utility theory:

Let w be a person's *wealth*.

Let $\mathbf{x} = (x_1, p_1; ...; x_n, p_n)$ be a lottery (or gamble or risky prospect).

- Note: $\mathbf{x}$ yields *income* x_i with probability p_i.

Five Features/Issues of Expected Utility

(1) Integration: *EU* operates on *final wealth*.

So thus far we have been doing things a bit sloppily.

Proper way to use expected-utility theory:

Let w be a person's *wealth*.

Let $\mathbf{x} = (x_1, p_1; ...; x_n, p_n)$ be a lottery (or gamble or risky prospect).

- Note: $\mathbf{x}$ yields *income* x_i with probability p_i.

EU theory says evaluate prospect $\mathbf{x}$ according to:

$$EU(\mathbf{x}; w) \equiv \sum_{i=1}^{n} p_i \, u(w + x_i).$$

Example: Demand for Insurance

Suppose you have wealth $1000. (Sorry, poor students)

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery (1000, .9 ; 750, .1).

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery (1000, .9 ; 750, .1).

Insurance agent: "I will fully insure you if you pay me \$$p$."

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery $(\ 1000, .9\ ;\ 750, .1\ )$.

Insurance agent: “I will fully insure you if you pay me \$$p$.”

- Implication: If buy insurance, you face the lottery $(\ 1000 - p\ ,\ 1\ )$.

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery $(\ 1000, .9\ ;\ 750, .1\ )$.

Insurance agent: "I will fully insure you if you pay me \$$p$."

- Implication: If buy insurance, you face the lottery $(\ 1000 - p\ ,\ 1\ )$.

What is your willingness to pay for full insurance?

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery $(\ 1000, .9\ ;\ 750, .1\ )$.

Insurance agent: "I will fully insure you if you pay me \$$p$."

- Implication: If buy insurance, you face the lottery $(\ 1000 - p\ ,\ 1\ )$.

What is your willingness to pay for full insurance?

- Answer: It's the p^* such that

$$(\ 1000 - p^*, 1\ ) \sim (\ 1000, .9\ ;\ 750, .1\ )$$

Example: Demand for Insurance

Suppose you have wealth \$1000. (Sorry, poor students)

But there is a 10% chance that you will suffer a loss of \$250.

- Hence, you face the lottery $(\ 1000, .9\ ;\ 750, .1\ )$.

Insurance agent: "I will fully insure you if you pay me \$$p$."

- Implication: If buy insurance, you face the lottery $(\ 1000 - p\ ,\ 1\ )$.

What is your willingness to pay for full insurance?

- Answer: It's the p^* such that

$$(\ 1000 - p^*, 1\ ) \sim (\ 1000, .9\ ;\ 750, .1\ )$$

  or

$$u(1000 - p^*) = (.9)\ u(1000) + (.1)\ u(750).$$

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Definition: $\tilde{u}(x)$ *and* $u(x)$ *reflect the same preferences* if for any pair of lotteries $\mathbf{x}$ and $\mathbf{y}$

$$EU(\mathbf{x}) > EU(\mathbf{y}) \text{ if and only if } E\tilde{U}(\mathbf{x}) > E\tilde{U}(\mathbf{y}).$$

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Definition: $\tilde{u}(x)$ *and* $u(x)$ *reflect the same preferences* if for any pair of lotteries **x** and **y**

$$EU(\mathbf{x}) > EU(\mathbf{y}) \text{ if and only if } E\tilde{U}(\mathbf{x}) > E\tilde{U}(\mathbf{y}).$$

Theorem: $\tilde{u}(x)$ and $u(x)$ reflect the same preferences if and only if $\tilde{u}(x)$ is a positive affine transformation of $u(x)$.

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Definition: $\tilde{u}(x)$ *and* $u(x)$ *reflect the same preferences* if for any pair of lotteries $\mathbf{x}$ and $\mathbf{y}$

$$EU(\mathbf{x}) > EU(\mathbf{y}) \text{ if and only if } E\tilde{U}(\mathbf{x}) > E\tilde{U}(\mathbf{y}).$$

Theorem: $\tilde{u}(x)$ and $u(x)$ reflect the same preferences if and only if $\tilde{u}(x)$ is a positive affine transformation of $u(x)$.

(3) Can $u(x)$ be negative? Can $EU(\mathbf{x})$ be negative?

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Definition: $\tilde{u}(x)$ *and* $u(x)$ *reflect the same preferences* if for any pair of lotteries $\mathbf{x}$ and $\mathbf{y}$

$$EU(\mathbf{x}) > EU(\mathbf{y}) \text{ if and only if } E\tilde{U}(\mathbf{x}) > E\tilde{U}(\mathbf{y}).$$

Theorem: $\tilde{u}(x)$ and $u(x)$ reflect the same preferences if and only if $\tilde{u}(x)$ is a positive affine transformation of $u(x)$.

(3) Can $u(x)$ be negative? Can $EU(\mathbf{x})$ be negative?

- Answer: Sure. Consider $b = -1000000$.

Five Features/Issues of Expected Utility

(2) The utility function is unique up to a positive affine transformation.

Definition: The function $\tilde{u}(x)$ is a *positive affine transformation* of the function $u(x)$ if $\tilde{u}(x) = a \cdot u(x) + b$ for some $a > 0$ and any b.

Definition: $\tilde{u}(x)$ *and* $u(x)$ *reflect the same preferences* if for any pair of lotteries $\mathbf{x}$ and $\mathbf{y}$

$$EU(\mathbf{x}) > EU(\mathbf{y}) \text{ if and only if } E\tilde{U}(\mathbf{x}) > E\tilde{U}(\mathbf{y}).$$

Theorem: $\tilde{u}(x)$ and $u(x)$ reflect the same preferences if and only if $\tilde{u}(x)$ is a positive affine transformation of $u(x)$.

(3) Can $u(x)$ be negative? Can $EU(\mathbf{x})$ be negative?

- Answer: Sure. Consider $b = -1000000$.
- Point: The level of utility doesn't matter, only the comparisons.

Five Features/Issues of Expected Utility

(4) Closely related: *EU* operates on *reduced lotteries*.

Five Features/Issues of Expected Utility

(4) Closely related: *EU* operates on *reduced lotteries*.

Example: If you face both lottery $\mathbf{x} = (100, 1/2; -90, 1/2)$ and $\mathbf{y} = (95, 1/2; -85, 1/2)$, you must first combine these into a reduced lottery before applying *EU*.

Five Features/Issues of Expected Utility

(4) Closely related: *EU* operates on *reduced lotteries*.

Example: If you face both lottery $\mathbf{x} = (100, 1/2; -90, 1/2)$ and $\mathbf{y} = (95, 1/2; -85, 1/2)$, you must first combine these into a reduced lottery before applying *EU*.

- If they are independent, the reduced lottery is

$$(\ 195, 1/4 \ ; \ 15, 1/4 \ ; \ 5, 1/4 \ ; \ -175, 1/4 \ ).$$

Five Features/Issues of Expected Utility

(4) Closely related: *EU* operates on *reduced lotteries*.

Example: If you face both lottery $\mathbf{x} = (100, 1/2; -90, 1/2)$ and $\mathbf{y} = (95, 1/2; -85, 1/2)$, you must first combine these into a reduced lottery before applying *EU*.

- If they are independent, the reduced lottery is

$$(\ 195, 1/4\ ;\ 15, 1/4\ ;\ 5, 1/4\ ;\ \ -175, 1/4\ ).$$

- If they have perfect positive correlation, the reduced lottery is

$$(\ 195, 1/2\ ;\ \ -175, 1/2\ ).$$

Five Features/Issues of Expected Utility

(4) Closely related: *EU* operates on *reduced lotteries*.

Example: If you face both lottery $\mathbf{x} = (100, 1/2; -90, 1/2)$ and $\mathbf{y} = (95, 1/2; -85, 1/2)$, you must first combine these into a reduced lottery before applying *EU*.

- If they are independent, the reduced lottery is

$$(\; 195, 1/4 \; ; \; 15, 1/4 \; ; \; 5, 1/4 \; ; \; -175, 1/4 \;).$$

- If they have perfect positive correlation, the reduced lottery is

$$(\; 195, 1/2 \; ; \; -175, 1/2 \;).$$

- If they have perfect negative correlation, the reduced lottery is

$$(\; 15, 1/2 \; ; \; 5, 1/2 \;).$$

Five Features/Issues of Expected Utility

(5) There is a natural mathematical way to describe the degree of risk aversion. (At least, it is natural in many ways, though opaque.)

Five Features/Issues of Expected Utility

(5) There is a natural mathematical way to describe the degree of risk aversion. (At least, it is natural in many ways, though opaque.)

Goal: We'd like to make statements of the form, "If person *A* is more risk-averse than person *B*, then *A* is less likely to accept a particular gamble."

Five Features/Issues of Expected Utility

(5) There is a natural mathematical way to describe the degree of risk aversion. (At least, it is natural in many ways, though opaque.)

Goal: We'd like to make statements of the form, "If person A is more risk-averse than person B, then A is less likely to accept a particular gamble."

Question: How to formalize "more risk averse"?

Definition: The *coefficient of risk aversion* at wealth w is

$$r(w) = \frac{-u''(w)}{u'(w)}.$$

Suppose w_A is person A's wealth and w_B is person B's wealth. If $r(w_A) > r(w_B)$ then person A is less likely than person B to accept a risky gamble.

Five Features/Issues of Expected Utility

Note: A person's risk aversion depends on both her utility function and her wealth.

$$r(w) = \frac{-u''(w)}{u'(w)}.$$

How is one's willingness to accept a gamble likely to depend on my wealth?

Five Features/Issues of Expected Utility

Note: A person's risk aversion depends on both her utility function and her wealth.

$$r(w) = \frac{-u''(w)}{u'(w)}.$$

How is one's willingness to accept a gamble likely to depend on my wealth?

- Economists often assume risk aversion decreases with wealth:

$$\uparrow w \Longrightarrow \downarrow r(w).$$

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0, \ \rho \neq 1.$$

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.
- Note that $r(w) = \rho/w$, hence $\uparrow \rho \Longrightarrow \uparrow r(w)$ and $\uparrow w \Longrightarrow \downarrow r(w)$.

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.
- Note that $r(w) = \rho/w$, hence $\uparrow \rho \Longrightarrow \uparrow r(w)$ and $\uparrow w \Longrightarrow \downarrow r(w)$.
- Person feels the same about the following gambles:

gamble $(\$11, 1/2; -\$10, 1/2)$ from wealth $\$1000$

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.
- Note that $r(w) = \rho/w$, hence $\uparrow \rho \Longrightarrow \uparrow r(w)$ and $\uparrow w \Longrightarrow \downarrow r(w)$.
- Person feels the same about the following gambles:

$$\text{gamble } (\$11, 1/2; -\$10, 1/2) \text{ from wealth } \$1000$$
$$\text{gamble } (\$110, 1/2; -\$100, 1/2) \text{ from wealth } \$10000$$

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.
- Note that $r(w) = \rho/w$, hence $\uparrow \rho \Longrightarrow \uparrow r(w)$ and $\uparrow w \Longrightarrow \downarrow r(w)$.
- Person feels the same about the following gambles:

$$\text{gamble } (\$11, 1/2; -\$10, 1/2) \text{ from wealth } \$1000$$
$$\text{gamble } (\$110, 1/2; -\$100, 1/2) \text{ from wealth } \$10000$$
$$\text{gamble } (\$1100, 1/2; -\$1000, 1/2) \text{ from wealth } \$100000$$

An Important Functional Form

Definition: The *CRRA utility function* is

$$u(w) = \frac{(w)^{1-\rho}}{1-\rho} \qquad \text{for } \rho \geq 0,\ \rho \neq 1.$$

Some features:

- For $\rho = 0$, equivalent to $u(w) = w$, or risk neutrality.
- For $\rho = 1$, we define $u(w) = \ln w$ — the natural continuation: $\ln w$ is the limit of $\frac{w^{1-\rho}-1}{1-\rho}$ as $\rho \to 1$, and subtracting the constant $\frac{1}{1-\rho}$ is just an affine shift, which changes nothing.
- Note that $r(w) = \rho/w$, hence $\uparrow \rho \Longrightarrow \uparrow r(w)$ and $\uparrow w \Longrightarrow \downarrow r(w)$.
- Person feels the same about the following gambles:

$$\text{gamble } (\$11, 1/2; -\$10, 1/2) \text{ from wealth } \$1000$$
$$\text{gamble } (\$110, 1/2; -\$100, 1/2) \text{ from wealth } \$10000$$
$$\text{gamble } (\$1100, 1/2; -\$1000, 1/2) \text{ from wealth } \$100000$$

Brief Calculus Refresher

Slope

- Recall: basic idea of a slope is the **rise** (change in y) over the **run** (change in x)

Slope

- Recall: basic idea of a slope is the **rise** (change in y) over the **run** (change in x)
- For straight lines (linear functions), the slope is the same everywhere:

$$y = f(x) = mx + b \quad \rightarrow \quad m \text{ is the slope}$$

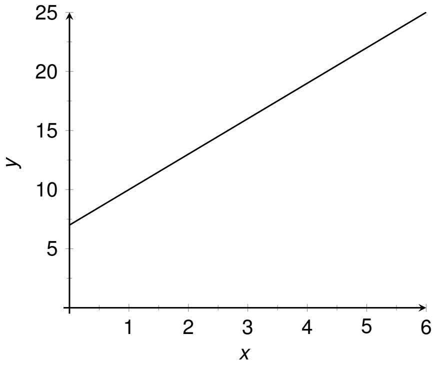

Slope

- For any two points on the graph, the slope of the line between them is

$$\text{Slope} = \frac{\text{rise}}{\text{run}} = \frac{f(x_2)-f(x_1)}{x_2-x_1}$$

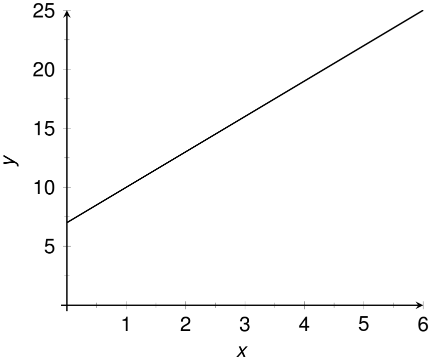

Slope

- For any two points on the graph, the slope of the line between them is

$$\text{Slope} = \frac{\text{rise}}{\text{run}} = \frac{f(x_2) - f(x_1)}{x_2 - x_1}$$

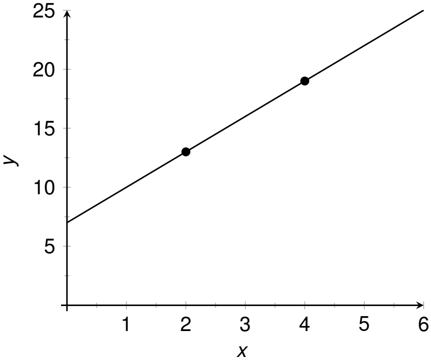

Slope

- For any two points on the graph, the slope of the line between them is

$$\text{Slope} = \frac{\text{rise}}{\text{run}} = \frac{f(x_2) - f(x_1)}{x_2 - x_1}$$

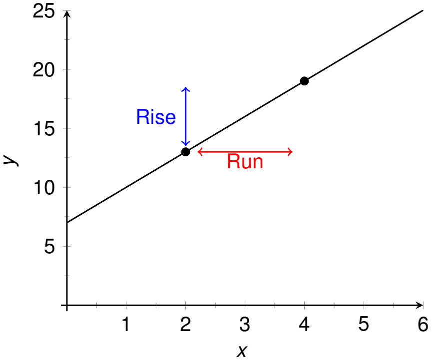

Slope and Curvature

- If the function is not linear, i.e., it has some curvature, taking the slope of a line between two points only *approximates* the slope

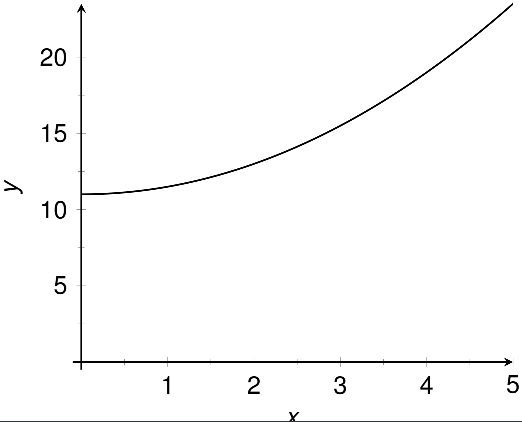

Slope and Curvature

- If the function is not linear, i.e., it has some curvature, taking the slope of a line between two points only *approximates* the slope

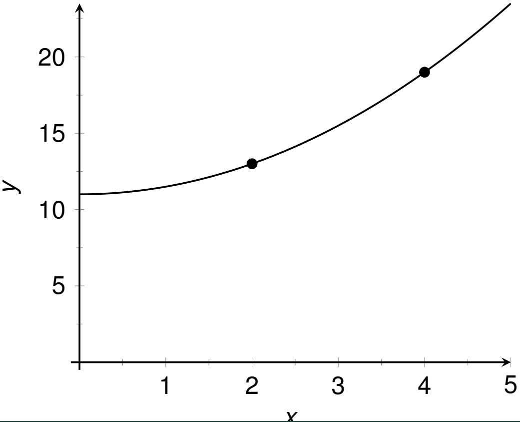

Slope and Curvature

- If the function is not linear, i.e., it has some curvature, taking the slope of a line between two points only *approximates* the slope

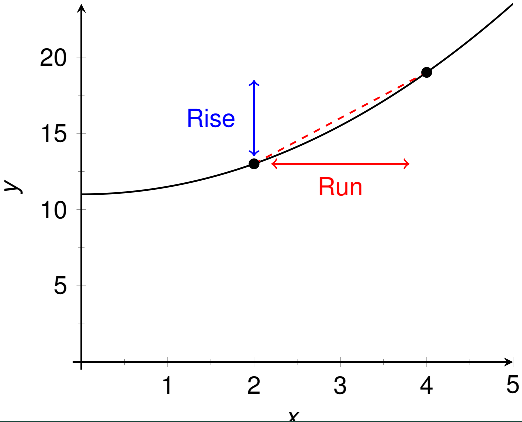

Slope and Curvature

- If the function is not linear, i.e., it has some curvature, taking the slope of a line between two points only *approximates* the slope
- Taking **derivatives** gives us a more precise measure of the slope

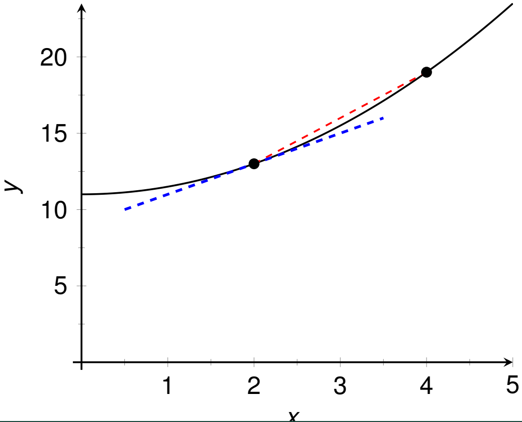

Derivatives

- **Idea:** to obtain a more precise approximation of the slope at $x_1 = x$, calculate the slope of the line between $(x_1, f(x_1))$ and $(x_2, f(x_2))$ as x_2 gets arbitrarily close to x_1
  - In other words, make the run $(x_2 - x_1)$ very small

Derivatives

- **Idea:** to obtain a more precise approximation of the slope at $x_1 = x$, calculate the slope of the line between $(x_1, f(x_1))$ and $(x_2, f(x_2))$ as x_2 gets arbitrarily close to x_1
  - In other words, make the run $(x_2 - x_1)$ very small
- Bonus material: Formally, the derivative is the limit of the slope obtained by using $x_1 = x$ and $x_2 = x + h$ as h gets closer and closer to 0

Derivatives

- **Idea:** to obtain a more precise approximation of the slope at $x_1 = x$, calculate the slope of the line between $(x_1, f(x_1))$ and $(x_2, f(x_2))$ as x_2 gets arbitrarily close to x_1
  - In other words, make the run $(x_2 - x_1)$ very small
- Bonus material: Formally, the derivative is the limit of the slope obtained by using $x_1 = x$ and $x_2 = x + h$ as h gets closer and closer to 0

$$\text{derivative} = \frac{df(x)}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{(x+h) - x}$$

Derivatives

- **Idea:** to obtain a more precise approximation of the slope at $x_1 = x$, calculate the slope of the line between $(x_1, f(x_1))$ and $(x_2, f(x_2))$ as x_2 gets arbitrarily close to x_1
  - In other words, make the run $(x_2 - x_1)$ very small
- Bonus material: Formally, the derivative is the limit of the slope obtained by using $x_1 = x$ and $x_2 = x + h$ as h gets closer and closer to 0

$$\text{derivative} = \frac{df(x)}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{(x+h) - x}$$

- Note: we can think of a derivative as the "marginal value" of a function
  - It tells us how much y would change if x were to increase by a small amount

Computing Marginal Values of Functions

- Rules of calculus can be applied to any function to obtain its derivative (or "marginal value") at any given point
- If you've not seen calculus before, there are a few useful rules to remember

Computing Marginal Values of Functions

- Rules of calculus can be applied to any function to obtain its derivative (or "marginal value") at any given point
- If you've not seen calculus before, there are a few useful rules to remember
- First, some notation: let $y = f(x)$ be a function
- The "derivative of y with respect to x" is denoted by

$$\frac{dy}{dx}$$

Computing Marginal Values of Functions

- Rules of calculus can be applied to any function to obtain its derivative (or "marginal value") at any given point
- If you've not seen calculus before, there are a few useful rules to remember
- First, some notation: let $y = f(x)$ be a function
- The "derivative of y with respect to x" is denoted by

$$\frac{dy}{dx}$$

  - Other way to word this: derivative of f with respect to x, denoted by $\frac{df(x)}{dx}$
  - Yet another way to write this: $f'(x)$

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$
- Example: $y = -100 \Rightarrow \frac{dy}{dx} = 0$

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$
- Example: $y = -100 \Rightarrow \frac{dy}{dx} = 0$
- Why?

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$
- Example: $y = -100 \Rightarrow \frac{dy}{dx} = 0$
- Why?
- Graphically: the graph of $y = m$ is a horizontal line $\Rightarrow$

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$
- Example: $y = -100 \Rightarrow \frac{dy}{dx} = 0$
- Why?
- Graphically: the graph of $y = m$ is a horizontal line $\Rightarrow$ it's neither increasing or decreasing $\Rightarrow$ derivative is zero

Derivative Rule 1: Constant Functions

- If y is a constant function of x, i.e. $y = b$ (a constant), then

$$\frac{dy}{dx} = 0.$$

- Example: $y = 42 \Rightarrow \frac{dy}{dx} = 0$
- Example: $y = -100 \Rightarrow \frac{dy}{dx} = 0$
- Why?
- Graphically: the graph of $y = m$ is a horizontal line $\Rightarrow$ it's neither increasing or decreasing $\Rightarrow$ derivative is zero
- Put differently, if $y = m$, then y doesn't depend on x at all
  - Thus, a small change in x doesn't change y at all $\Rightarrow$ derivative is zero!

Derivative Rule 2: Linear Functions (i.e., constant tim

- If y is a linear function of x, i.e. $y = mx + b$, then

$$\frac{dy}{dx} = m.$$

Derivative Rule 2: Linear Functions (i.e., constant tim

- If y is a linear function of x, i.e. $y = mx + b$, then

$$\frac{dy}{dx} = m.$$

- Example: $y = 10x + 2 \Rightarrow \frac{dy}{dx} =$

Derivative Rule 2: Linear Functions (i.e., constant tim

- If y is a linear function of x, i.e. $y = mx + b$, then

$$\frac{dy}{dx} = m.$$

- Example: $y = 10x + 2 \Rightarrow \frac{dy}{dx} = 10$

Derivative Rule 2: Linear Functions (i.e., constant tim

- If y is a linear function of x, i.e. $y = mx + b$, then

$$\frac{dy}{dx} = m.$$

- Example: $y = 10x + 2 \Rightarrow \frac{dy}{dx} = 10$
- Note that the function $y = mx + b$ is just a straight line with slope m
  - b is the y-intercept, which just shifts the graph up and down but doesn't affect its slope

Derivative Rule 2: Linear Functions (i.e., constant tim

- If y is a linear function of x, i.e. $y = mx + b$, then

$$\frac{dy}{dx} = m.$$

- Example: $y = 10x + 2 \Rightarrow \frac{dy}{dx} = 10$
- Note that the function $y = mx + b$ is just a straight line with slope m
  - b is the y-intercept, which just shifts the graph up and down but doesn't affect its slope
- The fact that dy/dx is a constant means that the slope is never-changing

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent
- Example: $y = 5x^2 \Rightarrow \frac{dy}{dx} =$

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent
- Example: $y = 5x^2 \Rightarrow \frac{dy}{dx} = 5 \cdot 2 \cdot x^{2-1} = 10x$

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent
- Example: $y = 5x^2 \Rightarrow \frac{dy}{dx} = 5 \cdot 2 \cdot x^{2-1} = 10x$
- Practice:
  - $y = 40x^3 \Rightarrow$

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent
- Example: $y = 5x^2 \Rightarrow \frac{dy}{dx} = 5 \cdot 2 \cdot x^{2-1} = 10x$
- Practice:
  - $y = 40x^3 \Rightarrow$
  - $y = \sqrt{x} = x^{\frac{1}{2}} \Rightarrow$

Derivative Rule 3: Functions that have variables raise

- If y is a power function of x, i.e. $y = ax^n$ where a and n are constants, then

$$\frac{dy}{dx} = anx^{n-1}.$$

- Simple procedure: Multiply the function by its exponent, and then subtract 1 from the original exponent
- Example: $y = 5x^2 \Rightarrow \frac{dy}{dx} = 5 \cdot 2 \cdot x^{2-1} = 10x$
- Practice:
  - $y = 40x^3 \Rightarrow$
  - $y = \sqrt{x} = x^{\frac{1}{2}} \Rightarrow$
  - $y = 5/x^3 = 5x^{-3} \Rightarrow$

Derivative Rule 4: Functions that have sums of variab

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.

Derivative Rule 4: Functions that have sums of variables

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.
- If $y = f(x) + g(x)$, then

$$\frac{dy}{dx} = \frac{df}{dx} + \frac{dg}{dx}.$$

Derivative Rule 4: Functions that have sums of varial

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.
- If $y = f(x) + g(x)$, then
$$\frac{dy}{dx} = \frac{df}{dx} + \frac{dg}{dx}.$$
- Example: $y = 3x^2 - 7x + 10 \Rightarrow \frac{dy}{dx} =$

Derivative Rule 4: Functions that have sums of varial

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.
- If $y = f(x) + g(x)$, then
$$\frac{dy}{dx} = \frac{df}{dx} + \frac{dg}{dx}.$$
- Example: $y = 3x^2 - 7x + 10 \Rightarrow \frac{dy}{dx} = 6x - 7$
- Practice:
  - $y = 9x + 163 \Rightarrow$

Derivative Rule 4: Functions that have sums of variables

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.
- If $y = f(x) + g(x)$, then
$$\frac{dy}{dx} = \frac{df}{dx} + \frac{dg}{dx}.$$
- Example: $y = 3x^2 - 7x + 10 \Rightarrow \frac{dy}{dx} = 6x - 7$
- Practice:
  - $y = 9x + 163 \Rightarrow$
  - $y = 0.5x^3 + 0.1x^2 - 9 \Rightarrow$

Derivative Rule 4: Functions that have sums of varial

- If the overall function is the sum of many functions: Simply split up the function, and take the derivatives of the individual bits. Then add them back together.
- If $y = f(x) + g(x)$, then
$$\frac{dy}{dx} = \frac{df}{dx} + \frac{dg}{dx}.$$
- Example: $y = 3x^2 - 7x + 10 \Rightarrow \frac{dy}{dx} = 6x - 7$
- Practice:
  - $y = 9x + 163 \Rightarrow$
  - $y = 0.5x^3 + 0.1x^2 - 9 \Rightarrow$
  - $y = -40x + 5x^{-2} \Rightarrow$

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then

$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then

$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$

- Example: $u(x) = (5 + 2\alpha x)^3 \Rightarrow \frac{du}{dx} =$

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then
$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$
- Example: $u(x) = (5 + 2\alpha x)^3 \Rightarrow \frac{du}{dx} = 3(5 + 2x)^2 \cdot (2)$
- Practice:
  - $y = \sqrt{3x + x^3} \Rightarrow$

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then
$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$
- Example: $u(x) = (5 + 2\alpha x)^3 \Rightarrow \frac{du}{dx} = 3(5 + 2x)^2 \cdot (2)$
- Practice:
  - $y = \sqrt{3x + x^3} \Rightarrow$
  - $y = \ln(2 + 3x^2) \Rightarrow$

Derivative Rule 5: Chain Rule (composition of functio

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then
$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$
- Example: $u(x) = (5 + 2\alpha x)^3 \Rightarrow \frac{du}{dx} = 3(5 + 2x)^2 \cdot (2)$
- Practice:
  - $y = \sqrt{3x + x^3} \Rightarrow$
  - $y = \ln(2 + 3x^2) \Rightarrow$
- Note: There are a few other derivative rules worth knowing (product rule, quotient rule, etc.) but we can address those when needed.

Derivative Rule 5: Chain Rule (composition of functic

- If the overall function is the composition of two functions, then take the derivative of the outside function evaluated at the inside one, and multiply by the derivative of the inside.
- If $y = f(g(x))$, then
$$\frac{dy}{dx} = \frac{df(g(x))}{dx} \cdot \frac{dg}{dx}.$$
- Example: $u(x) = (5 + 2\alpha x)^3 \Rightarrow \frac{du}{dx} = 3(5 + 2x)^2 \cdot (2)$
- Practice:
  - $y = \sqrt{3x + x^3} \Rightarrow$
  - $y = \ln(2 + 3x^2) \Rightarrow$
- Note: There are a few other derivative rules worth knowing (product rule, quotient rule, etc.) but we can address those when needed. Derivative tables are useful too.

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0

Unconstrained Maximization

To find *x* value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow$$

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x =$

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x = 3$

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x = 3$
  - But is function maximized at $x = 3$ or minimized?

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x = 3$
  - But is function maximized at $x = 3$ or minimized?
  - Second derivative: $f''(x) = -2$

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x = 3$
  - But is function maximized at $x = 3$ or minimized?
  - Second derivative: $f''(x) = -2$
    - If second derivative is positive, then function convex
    - If it is negative, function is concave

Unconstrained Maximization

To find x value that maximizes a function:

1. Take derivative and find places with slopes equal to 0
2. To tell if these places represent maximums or minimums: determine if function is concave/convex at those places (2nd Derivative Test)

- **Example:**

$$f(x) = 6x - x^2 \quad \Rightarrow \quad f'(x) = 6 - 2x$$

  - $f'(x)$ is zero at $x = 3$
  - But is function maximized at $x = 3$ or minimized?
  - Second derivative: $f''(x) = -2$
    - If second derivative is positive, then function convex
    - If it is negative, function is concave
  - We are in the concave case in this example, and hence function is maximized at $x = 3$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs $\Rightarrow \pi = Q(100 - Q) - 10Q$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs $\Rightarrow \pi = Q(100 - Q) - 10Q$
- Take derivative of π and set it equal to 0

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs $\Rightarrow \pi = Q(100 - Q) - 10Q$
- Take derivative of π and set it equal to 0

  $\Rightarrow \frac{d\pi}{dQ} = 90 - 2Q = 0$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs $\Rightarrow \pi = Q(100 - Q) - 10Q$
- Take derivative of π and set it equal to 0

  $\Rightarrow \frac{d\pi}{dQ} = 90 - 2Q = 0 \Rightarrow Q^* = 45$

Unconstrained Maximization: Example

- **Goal:** find a monopoly's profit-maximizing quantity to produce
- Assumptions:
  - Demand curve: $P = 100 - Q$
  - Cost of producing each unit: $c = 10$
- Revenue from selling Q units: $Q \cdot P = Q \cdot (100 - Q)$
- Cost from producing Q units: $10 \times Q$

  $\Rightarrow$ Profit $=$ Revenue $-$ Costs $\Rightarrow \pi = Q(100 - Q) - 10Q$
- Take derivative of π and set it equal to 0

  $\Rightarrow \frac{d\pi}{dQ} = 90 - 2Q = 0 \Rightarrow Q^* = 45$
- We will do a more relevant example in Tuesday's class (similar to question on PS1)

Next Class

Today: The standard model of choice under uncertainty

Next class: Wrap up the standard model (applications), then the evidence against it begins

Remember:

- **Read the syllabus!**
- Problem Set 0 is posted (optional review of optimization) – no need to turn this in
- Problem Set 1 is posted – due Tuesday, September 15